

**REST AVAILABLE COPY**

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
30 August 2001 (30.08.2001)

PCT

(10) International Publication Number  
**WO 01/63479 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**Drive, Park Ridge, NJ 07656-1103 (US). **DONOGHUE, Karen**; 122 Lake Street, Arlington, MA 02474 (US).

(21) International Application Number: PCT/US01/40173

(74) Agent: **PRAHL, Eric, L.**; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).

(22) International Filing Date: 22 February 2001 (22.02.2001)

(81) Designated States (national): CA, JP.

(25) Filing Language: English

(84) Designated States (regional): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

(26) Publication Language: English

**Published:**


- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

(30) Priority Data:  
60/183,971 22 February 2000 (22.02.2000) US  
60/201,839 3 May 2000 (03.05.2000) US

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(71) Applicant: **METACARTA, INC.** [US/US]; P.O. Box 397207, Cambridge, MA 02139 (US).(72) Inventors: **FRANK, John, R.**; P.O. Box 397207, Cambridge, MA 02139 (US). **RAUCH, Erik, M.**; 40 Circle

(54) Title: SPATIALLY CODING AND DISPLAYING INFORMATION



(57) Abstract: A computer system (20) includes storage system (22) which contains information in the form of documents, along with a spatial information about the documents. The computer system (20) also includes subsystems for data collection (30), data analysis (40), search (50), data presentation (60), and portal services (70). The computer system (20) further includes a map interface (80). Through the map interface (80), users can query the storage (22) and view a representation of the query results arranged on a map.

WO 01/63479 A1

## Spatially Coding and Displaying Information

Under 35 U.S.C. §119(e)(1), this application claims benefit of prior U.S. Provisional Applications No. 60/183,971, entitled "Metacarta: Map-based Information Search Engine and Catalog," filed February 22, 2000; and No. 60/201,839, entitled "Method and System for Associating Information with Physical Objects and Locations; and Methods of Expanding a Database," filed May 3, 2000, both of which are incorporated herein by reference.

### TECHNICAL FIELD

This invention relates to computer systems, and more particularly to spatial databases, document databases, search engines, and data visualization.

10

### BACKGROUND

There are many tools available for organizing and accessing documents through different interfaces that help users find information. Some of these tools allow users to search for documents matching specific criteria, such as containing specified keywords. Some of these tools present information about geographic regions or spatial domains, such as driving directions presented on a map.

15

These tools are available on private computer systems and are sometimes made available over public networks, such as the Internet. Users can use these tools to gather information.

### SUMMARY OF THE INVENTION

20

In a computer system that presents a map interface to a user, the invention enables a user, among other things, to pose a query via the map interface and to be able to inspect a representation of the query results arranged on the map as icons. The map and the icons are responsive to further user actions, including changes to the scope of the map, changes to the terms of the query, or closer examination of a subset of the results.

25

The targets of the query are documents. Examples of documents include text-based computer files, as well as files that are partially text-based, files containing spatial information, and computer entities that can be accessed via a document-like interface.

Documents can contain other documents and may have other interfaces besides their document-like interfaces. Every document has an address. In the case of world wide web documents, this address is commonly a URL.

The documents exist on computer systems arrayed across a computer network, such as a private network or the Internet. The documents may be hyperlinked, that is, may contain references (hyperlinks) to an address of another document. Copies of the documents may be stored in the page repository.

A spatial recognizer process examines documents for spatial information content. When the spatial recognizer determines that a document has spatial information content, the document is added to a spatial document collection.

A document ranking process assigns a spatial relevance score to each document in the spatial document collection. The spatial relevance score is a measure of the degree to which the document relates to the spatial location mentioned in its spatial information content. In cases where the document has more than one instance of spatial information content, the document is scored against each instance.

The spatial-keyword document indexer examines each document in the spatial document collection and represents it in a spatial-keyword document index data structure. The spatial-keyword document indexer indexes a document both by keywords and by at least one instance of spatial information content. The spatial-keyword document index enables unusually fast responses by the computer system to queries that combine spatial criteria with keyword criteria.

The crawler extends the collection of known documents by examining the hyperlinks contained in the known documents. When a hyperlink references a previously unknown document, the crawler adds the unknown document to the collection of known documents and examines them, in turn, for new hyperlinks to follow.

The crawler may prioritize the hyperlinks it follows based in part on spatial relevance scores.

The computer system includes a metasearcher process for initializing the collection of known documents. This initializing step is known as bootstrapping and is known in the art. The metasearcher queries predetermined search engines known to store information about other computer systems and document sources, such as search engine web sites on the

Internet. The human administrators of the metasearcher provide it with a collection of known spatial locations. The metasearcher formulates queries based on these spatial locations and directs the queries to the search engines. After each query, the results are compared to the collection of known documents and are added if new.

5 However, it is common for search engines to cap the maximum number of results returnable to a single query. The metasearcher is able to respond to a results cap by issuing follow-on queries which are progressively more spatially focused. An example of a progressively more spatially focused series might be "New York state," "New York, NY," "Times Square, New York, NY," etc. By progressively narrowing the scope of its queries, 10 the metasearcher reduces the number of results until the results number fits within the cap. The progressive spatial focus produces information more closely matched to a specific spatial location, as well as a more exhaustive sample of the results available from a given search engine. At the same time, the generality of early queries casts a net as broad as possible, so as not to miss any results. As a result, the documents found by the metasearcher form a 15 diverse yet highly spatially-qualified sample for the crawler to start from.

In general, in one aspect, the invention is an interface program stored on a computer-readable medium for causing a computer system with a display device to perform a set of functions. The functions are accepting search criteria from a user including a free text entry query and a domain identifier identifying a domain; in response to accepting the search criteria from the user, retrieving a plurality of record identifiers each of which identifies a corresponding record which: (1) has associated therewith a location identifier that locates it at a specific location within the domain identified by the domain identifier; and (2) contains information that is responsive to the free text entry query; displaying a representation of the domain on the display device; and displaying on the display device a plurality of icons as representations of the records identified by the plurality of record identifiers. For each of the plurality of record identifiers, a corresponding one of the plurality of icons is displayed within the representation of the domain that is being displayed on the display device. The corresponding icon for each of the plurality of record identifiers is positioned within the representation of the domain at a coordinate within the domain that corresponds to the 20 location identifier for the corresponding record.

25

30

Preferred embodiments include one or more of the following features. The domain is a geographical region and the representation is a multi-dimensional map of the geographical region. More specifically, the representation is a two-dimensional map of the geographical region. The step of accepting input further includes accepting a designation by the user of a designated category, wherein each of the records corresponding to the plurality of retrieved record identifiers also includes information that falls within the designated category. The step of accepting the designation by the user of a category includes presenting to the user a list of predefined categories and accepting as the designated category a selection by the user from that list. The interface program also is for causing the computer to perform the further functions of, after displaying the corresponding icon for each of the plurality of record identifiers, accepting further search criteria from the user. The further search criteria are selected from the group of search criteria types consisting of a domain identifier input type, a free text entry query input type, and a category type. It also causes the computer, in response to accepting the further search criteria from the user, to perform the functions of: (1) retrieving a subset of the plurality of record identifiers, wherein the subset of the plurality of record identifiers identifies all record identifiers among the plurality of record identifiers that fall within the further search criteria; (2) displaying a two-dimensional map of a revised geographical region on the display device that is responsive to the further search criteria; (3) for each of the record identifiers of the subset of plurality of record identifiers, displaying a corresponding icon within the displayed map, wherein the corresponding icon for each of the record identifiers of the subset of the plurality of record identifiers is positioned within the displayed map at a coordinate that corresponds to the location identifier for the corresponding record; and (4) storing as a filter the first-mentioned search criteria in combination with the further search criteria, wherein the stored filter is retrievable for use by the user in specifying a future search through the interface. The first-mentioned search criteria in combination with the further search criteria is an ordered sequence of inputs and the stored filter is the ordered sequence of inputs preserving the order of the sequence of inputs.

Preferred embodiments may also include one or more of the following features. The interface program also causes the computer to perform the further functions of: presenting to the user via the display device a map; and enabling the user to input the domain identifier as

part of the search criteria by interacting with the displayed map. The plurality of icons include an icon of a first icon class and an icon of a second icon class; and the icon of the first icon class has first visual characteristics and the icon of the second icon class has second visual characteristics that are different from the visual characteristics associated with the first icon class. At least some of the records identified by the plurality of record identifiers are of a first type and at least some of the other records identified by the plurality of record identifiers are of a second type and the records of the first type are displayed using the icon of the first icon class and records of the second type are displayed using the icon of the second icon class. At least one of the icons of the plurality of icons represents multiple of the records identified by the plurality of record identifiers, wherein each of the multiple of the records having a location identifier that locates that record within a neighborhood about a central location.

Also in preferred embodiments, the interface program causes the computer to perform the further functions of: accepting a change of scale request from the user; in response to accepting the change of scale request, consolidating at least some of the plurality of icons with each other to form a second plurality of icons that is fewer in number than the number of icons in the first-mentioned plurality of icons; and in response to accepting the change of scale request, redisplaying the domain using a decreased scale and also using the second plurality of icons to identify the locations of the records identified by the plurality of record identifiers. In addition, it also causes the computer to perform the further functions of: accepting from the user a specification of an electronic note which has an associated location within the map; and displaying a sticky-note icon on the map at position that corresponds to the associated location. The electronic note comprises a web page with its own externally accessible address which enables people to electronically access its contents through that address.

In general, in another aspect, the invention is a method that performs the functions described above.

In general, in yet another aspect, the invention is a database system stored on a computer-readable medium for causing a computer system to perform the functions of: accepting search criteria including at least one of: (1) text; (2) a domain identifier identifying a domain; and (3) a filter identifier identifying a filter; and retrieving a plurality of record

identifiers each of which identifies a corresponding record which has associated therewith the text, domain identifier, or layer identifier of the search criteria, where the retrieving is performed with a spatial-keyword document index.

In general, in still another aspect, the invention is a method of retrieving a plurality of record identifiers each of which identifies a corresponding record which has associated therewith at least one of a text, a domain identifier, and a layer identifier as specified by search criteria, wherein such retrieving is performed with a spatial-keyword document index.

Preferred embodiments include one or more of the following features. The spatial-keyword document index includes a spatial index tree extended to reference documents and a plurality of trees with the same structure as the spatial index tree but trimmed for particular lexicon entries and filters. The plurality of record identifiers are retrieved in a spatial-keyword document index tree and the branching structure of the tree is analyzed to identify geographic phenomena. A geographic phenomena is a spatial-keyword document index tree branching structure in which more than a predetermined fraction of the branches share more than predetermined number of parent nodes.

In general, in still yet another aspect, the invention is a program stored on a computer-readable medium for causing a computer system to perform the functions of: loading documents referred to by document addresses; parsing those documents for more document addresses to load; and parsing those documents for possible spatial identifiers.

Preferred embodiments may include one or more of the following features. The program also causes the computer system to perform the further function of analyzing the possible spatial identifiers to determine a location in a domain. A portion of the document addresses are gathered by a metasearcher process, which queries other computer systems using text that refers to spatial domains. The program also causes the computer system to perform the further function of computing a relevance score for each of the possible spatial identifiers found in each document. The relevance score includes one or more of: (1) the possible spatial identifier's position in the document; (2) the number of other possible spatial identifiers in the document; (3) whether the possible spatial identifier is in a sentence or is free standing; and (4) the formatted emphasis of the characters in the possible spatial identifier. The program also causes the computer system to perform the further function of sorting the document addresses by relevance score before loading.

In general in still another aspect, the invention is a program stored on a computer-readable medium for causing a computer system to perform the functions of: loading documents referred to by document addresses; parsing those documents for more document addresses to load; parsing those documents for possible spatial identifiers; and analyzing the possible spatial identifiers to determine a location in a domain.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

## DESCRIPTION OF DRAWINGS

FIG. 1 schematically shows an overall arrangement of a computer system according to an embodiment of the invention;

FIG. 2 schematically represents an arrangement of controls on a map interface according to an embodiment of the invention;

FIG. 3 is an explanatory diagram of storage entities and entities in a data collection process;

FIG. 4. is an explanatory diagram of entities in a data analysis process;

FIG. 5. is an explanatory diagram of entities in a search process;

FIG. 6. is an explanatory diagram of steps in a process for building a spatial-keyword indexer; and

FIG. 7. is an explanatory diagram of steps in a spatial indexer process.

## DETAILED DESCRIPTION

In general, with reference to Fig. 1, the computer system 20 includes a storage 22 system which contains information in the form of documents, along with spatial information about the documents. The computer system 20 also includes subsystems for data collection 30, data analysis 40, search 50, data presentation 60, and portal services 70. The computer system 20 further includes a map interface 80 presented to a user through a variety of clients. Through the map interface 80, the user can query the storage 22 and can view a representation of the query results arranged on a map.

## DOCUMENTS

The targets of a user query are documents. Examples of documents include text-based computer files, as well as files that are partially text-based, non-text files, files containing spatial information, and computer entities that can be accessed via a document-like interface. Documents can contain other documents and may include other interfaces besides their document-like interfaces. Every document has an address. In the case of World Wide Web documents, this address is commonly a URL. As is the case with URL's, a portion of the address may include instructions or parameters that are passed to the computer server process that serves the document.

The documents exist on computer systems arrayed across a computer network, such as a private network or the Internet. The documents may be hyperlinked, that is, may contain an address of another document. Copies of the documents may be stored in the page repository 222 (Fig. 3).

## GENERAL USER INTERFACE

With reference to Fig. 2, the map interface 80 is presented to the user on a computing device having a user interface. The user interface may be graphical (GUI), voice-based, or text-only. Each feature of the GUI will be reproduced in a voice-based or text-only user interface, to the extent possible.

As is common in the art, the GUI includes a pointer symbol responsive to the user's manipulation of a pointing device such as a mouse, a touch-sensitive area, or a combination of directional buttons. The pointer symbol is superimposed on the GUI contents. The GUI is also responsive to a click event generated by the user. The click is usually associated with the user's manipulation of a button on or near the pointing device, but may be activated in other ways, depending on the computing device and its operating system. The client process receives click events and the position of the pointer symbol from the operating system of the computing device.

The map interface 80 includes a map 805. The map 805 is a representation, often in part, of at least one spatial domain. A spatial domain is any space with a location metric known to the spatial recognizer 48. In one embodiment, the surface of the Earth is a spatial domain under the 2-dimensional location metric of latitude and longitude – henceforth, the

"standard geographic domain." In another embodiment, the "GPS domain" is defined by a volume around the surface of the Earth under the 3-dimensional metric of GPS (global positioning satellite) data.

A metric on a space need not identify spatial point locations. A document may be  
5 identified as being near a spatial point location. For example, a document could be identified as "near exit 19 off I-80 in Pennsylvania." A document could refer to an extended region like Plum Island state park.

The map 805 uses a scale in representing the domain. The scale indicates what subset  
of the domain will be displayed in the map 805. There is usually a range of scales  
10 appropriate to a given domain. By choosing a smaller scale, the user can examine a smaller  
portion of the entire domain in exchange for more detail per unit.

Several domains may exist along one spatial continuum. For instance, in one  
embodiment, the map 805 initially displays a representation of the whole surface of the Earth  
under the standard geographic domain. The map 805 might then be changed to display only  
15 one continent. This map 805 still displays a portion of the domain, but the scale has changed.  
However, the scale might change to the point that the map 805 displays only a close-up of a  
concert hall. At that point, the map interface 80 can change the domain to that of the concert  
hall, where locations can be expressed by section, row, and seat number, for instance.  
Domains can intersect or overlap, therefore, along a continuum of scale.

20 When the domain has geographic meaning, the map 805 may include standard  
geographic map features such as streets and waterways. Data for geographic features is  
available from the U.S. Census Bureau, the U.S. Geographic Survey, and companies such as  
GDT, of 11 Lafayette Street, Lebanon, NH, or NavTech, of 10400 W. Higgins Road,  
Rosemont, IL. The map can include spatial landmark features appropriate only to certain  
25 scales within the same domain, such as street names that do not appear until the scale is  
sufficient to allow individual streets to be represented.

The map 805 need not represent a domain that physically exists; the map 805 may  
represent something that is itself a representation, such as a virtual layout of a planned  
housing development. Still more abstractly, the map 805 may represent entities in a spatial  
30 layout where the spatial dimensions do not correspond to physical spatial dimensions. For

instance, the domain may be a genealogical tree laid out on a plane, wherein one axis of the plane represents the linear progress of time.

A domain location is a location in the space that the map 805 represents. The domain location is usefully distinguished from the display location, which describes the placement of elements as displayed by the map 805. The measure of distance between domain locations depends on the domain, whereas the distance between display locations is measured in pixels of the computing device hosting the client 64.

A domain frame is the subset of the overall domain (possibly including the entire domain) displayed by a state of the map 805.

10

#### *CHANGING THE MAP VIEW*

The user can adjust the view displayed by the map 805 in several ways.

The user can change the scale of the map 805 by a click on the zoom bar 891. The zoom bar 891 visually represents a plurality of scales that the map interface 80 is capable of displaying in the map 805. The scales displayed by the zoom bar 891 in any given state may be a subset of the total scales the map interface 80 is capable of displaying in the map 805.

This subset may adapt to a change in state, including a change in scale. For instance, in a geographic context, the first state of the map interface 80 may display the entire globe in the map 805. In this first state, the zoom bar 891 may display scales ranging from global to street level, for instance. If the user narrows the scope of the displayed region to a street address corresponding to a concert hall, the zoom bar 891 might display scales ranging from seating sections to individual seats inside the hall.

The user can move the center of the map 805 by a click on the map border 892. The map border 892 surrounds the map 805.

By interacting with the map mode controls 830, the user can specify how the map interface 80 should respond to clicks on the map 805. The map mode controls 830 include controls for pan 832, zoom 834, and post a note 836. The pan 832 control and the zoom 834 control each have states including an "on" state and an "off" state. When the pan 832 control is in its "on" state, a click on the map 805 instructs the map interface 80 to re-center the map 805 around the location represented by the click. Likewise, when the zoom 834 control is in

its "on" state, a click on the map 805 instructs the map interface 80 to zoom the map 805 in around the location represented by the click. The post a note 836 control is described in the section on electronic notes, below.

The map interface 80 includes spatial criteria entry controls 806. The spatial criteria entry controls 806 include a data entry control 808, a submission control 809, and a prompt for spatial criteria 807. The prompt for spatial criteria 807 instructs the user as to the purpose of the data entry control 808. The prompt for spatial criteria 807 may include a static instruction or may be dynamically responsive to user interaction, such as movement of the pointer symbol over the data entry control 808. The prompt for spatial criteria 807 may include sound. The user invokes the submission control 809 to notify the client process 64 that the data in the data entry control 808 is complete.

Examples of spatial criteria include geographic measurements such as latitude, longitude, or altitude; postal address information; or, returning to the concert hall example, row and seat number. Spatial criteria also include criteria that are indirectly spatial, i.e., criteria that do not describe spatial attributes but specify an entity which does have spatial attributes. An example of indirectly spatial criteria is the tracking number of a package. The tracking number might not be spatially descriptive in itself, but the package at any given instant might have a last known location that could be represented on the map 805.

The map interface 80 includes keyword entry controls 801. The keyword entry controls 801 include a data entry control 803, a submission control 804, and a prompt for keywords 802. The prompt for keywords 802 instructs the user as to the purpose of the data entry control 803. As with the prompt for spatial criteria 807, the prompt for keywords 802 may include a static instruction or may be dynamically responsive to user interaction, and may include sound. The role of the submission control 804 within the keyword entry controls 801 is identical to the role of the submission control 809 within the spatial criteria entry controls 806.

Note that part or all of the spatial criteria entry controls 806 and the keyword entry controls 801 may use the same interface components. For instance, if the user enters text "shoes near to Cambridge, MA," the system may consider this both spatial criteria and keyword criteria.

Examples of keywords include any word of interest to the user, or simply a string pattern. The computer system 20 compares the data in the data entry control 803 against the contents of the documents in storage 22. There are no predetermined restrictions on the keywords that may match a document.

5       The user may enter any text he/she desires in the text entry tools. The computer system 20 will parse entries to get possible domain changing commands and keyword queries. Keyword queries can be of any form. The computer system 20 does not restrict the entries to predefined categories. Instead, the computer system 20 attempts to match the query text against text found in all documents in the corpus.

10       One way to match the query text is to split it into separate strings divided by white space, where white space is commonly defined in the art as tabs, spaces, carriage returns, and other characters generally referred to by the regular expression character "\s". Each of these separate strings can then be searched for in the documents.

15       The text contents of the documents can be similarly divided into separate strings divided by white space. Thus, if the text entered by the user match any strings in the document corpus, the computer system 20 can retrieve results.

20       This "free text entry query" allows much more versatile searching than searching by predetermined categories.

## 20            *ICONS*

The map interface 80 may include one or more icons 810 superimposed upon the map 805. Icons 810 need not be present in the map interface 80 when the client 64 is initially presented to the user. After the user has submitted a query, though, the map interface 80 may use icons 810 to represent documents in storage 22 that satisfy the query criteria to a degree determined by the search 50 process.

25       The display placement of an icon 810 represents a correlation between its documents and the corresponding domain location. Specifically, for a given icon 810 having a domain location, and for each document associated with the icon 810, the subsystem for data analysis 20 must have determined that the document relates to the domain location. The subsystem for data analysis 20 might determine such a relation from a user's inputting that location for

the document. Note that a document can relate to more than one domain location, and thus would be represented by more than one icon 810.

An individual icon 810 belongs to an icon class. Icons 810 of the same icon class share visual characteristics that may include shape, color, size, indexing scheme (Roman numerals versus letters, e.g.), or animated behavior. An icon face 818 is an interface element of the map interface 80 satisfying the requirements of an icon class. In one embodiment, the client process 64 runs on a computer equipped with a monitor having a pixel size of approximately 0.28 mm, which is approximately the industry standard for desktop computers at the time the invention was made. For this pixel size, typical icons would be 15 to 20 pixels in diameter.

Note that there may be more than one way to satisfy the requirements of an icon class, so an icon class may have more than one icon face 818. For an example, see icon subclasses, below.

Visual similarities conferred by icon class may be used to represent topical similarities among the documents the icon 810 represents. For instance, documents affiliated with restaurant menus might be represented by icons 810 sharing a fork-and-knife shape. The fork-and-knife shape would be a property of the icon class.

Different colors, shapes, tints, and animated motions of the icons 810 might represent different features of the documents represented by the icons 810.

A class of icons may share the same geometric shape but have different colors, or different shades of the same color. The different shades might represent the several different properties of the documents represented by the icons. Different properties of the documents include the time elapsed since the document was created, the time elapsed since the document was introduced to the system, a relevance measure of the document, the size of the document.

Another feature of the icon class is the icon subclass. Two icon classes may be subclasses of a third class if they share the characteristics affiliated with the third class but vary at least one other characteristic in a consistent and meaningful way. For instance, the icon class for restaurants might have subclasses for quality, as measured by a certain newspaper's restaurant reviews. All icons 810 in the icon subclasses for restaurant quality would have a fork-and-knife shape in common, but icons 810 would be colored green for

good reviews, red for poor reviews, or yellow for mixed-opinion reviews. They could even be divided in pie-chart fashion to show a percentage of each type of review. Thus, broad visual similarities can be used to imply broad topical similarities on one level, while visual sub-variation can be used to imply topical sub-variation on a second level. The icon legend 5 812 can inform the user of such conventions.

If icon class B is a subclass of icon class A, then icon class A is a parent class of icon class B.

Many domain locations have multiple documents referring to that location. To illustrate this to the user, the icon 810 used at that location in the map interface might be of a 1 0 different size, color, or shape from other icons 810. For example, an icon 810 may be made to appear "stacked" as if a few of the icons 810 were placed nearly on top of each other. For another example, the icon 810 might appear to have parts of different icons 810 spliced together.

In a preferred embodiment, different colored icons 810 represent different layers of documents; varying shapes represent varying numbers of documents; varying shades represent varying relevance numbers for the underlying documents. The relevance of a set of documents referring to a given domain location might be computed by averaging or summing the relevance of the individual documents.

An icon 810 may represent one location in a domain or several neighboring locations. 20 The number of locations depends upon several factors, including the scale of the map 805. When multiple icons 810 have display locations within a tolerance determined by the computer system 20, the map interface 80 consolidates icons 810 to increase visual clarity. Should the user change the scale of the map by zooming it, the map interface 80 recalculates whether to consolidate icons 810. The tolerance beyond which icon consolidation occurs 25 may vary. The primary factor in the decision to consolidate is whether icons 810 are overlapping. For many icons 810, a good test of overlap is whether the display locations are closer than two times the average diameter of the icon faces 818. Other factors in the decision to consolidate include visual characteristics of the icon faces 818, visual characteristics of the map 805, characteristics of the domain, characteristics of the 30 documents, and the number and variety of icons currently present in the display.

A consolidated icon 810 may represent multiple spatial domains. For instance, consider the standard geographic domain that includes Washington, DC, and another domain for Ford's Theater under the concert-hall seating metric. At certain display scales, documents that represent the Lincoln Memorial might be consolidated into the same icon as 5 documents describing the specific seat in Ford's Theater in which Lincoln was shot. In this example, the Lincoln Memorial document might be affiliated with the standard geographic domain. The Ford's Theater document might be affiliated with a domain specific to Ford's Theater, but in this example is may be displayed on the standard geographic because the entire Ford's Theater domain can be mapped onto a fairly small region, relative to the size of 10 the domain requested by the user.

An icon 810 may also represent multiple topical categories among its documents, regardless of whether the icon 810 is consolidated. In this case, the icon face 818 may be altered to reflect the multiplicity of topics.

The icon legend 812 is another element of the map interface 80. The icon legend 812 15 relates an icon 810 to the documents it represents. The icon legend 812 comprises a listing of documents. The listing may be grouped or ordered in a variety of ways.

Icons 810 are listed in the icon legend 812 according to an order compiled by the search 50 process.

A non-consolidated icon 810 represents a single display location. The order of its 20 documents as listed in the icon legend 812 is based on a relevance ranking compiled by the search 50 process. The relevance ranking scores each document against the user's query criteria.

A consolidated icon 810 may represent a plurality of domain locations. A consolidated icon 810 may represent a plurality of icon classes. The different icon classes 25 may entail different topical categories. The icon legend 812 may differentiate the document listings according to these topical categories: for instance, by grouping them by category; by adding a field to each entry in the list, specifying the category; or by adding a visual emphasis. The visual emphasis may include a change in typeface, a change in color, or the presence of an icon type affiliated with the category. Several effects can be combined, such 30 as grouping by category in combination with a variation in background color between adjacent groups.

## FILTERS

The map interface 80 includes two groups of controls for managing filters, a general filter display 850 and a user-specific filter display 860.

With reference to Fig. 3, a filter selects a subset of the corpus of documents in the page repository 222. Filters are defined recursively: a filter is a list of elements, where each element can be either a keyword string, a set of spatial criteria, a human-compiled list of documents, a domain frame, or another filter. The elements may be defined in a sequence allowing the user to select a collection of documents. The sequence of filters may be combined with the Boolean AND operator to produce an intersected document set that is the same for any order of the filters. Two sets of filters may be combined with the Boolean OR operator. When viewing a set of documents in a map 805, the user may change the map view to display a subset of this document set, which may be different than if the user performed the filtering operation after changing the map view. Thus, every user query defines a filter, because it contains either keywords, spatial criteria, a change to the domain frame, or several in combination. The initial state of the map interface 80 – even if the user has not yet interacted with it – defines a filter, since the map 805 has at least a domain frame associated with it. Similarly, because a non-empty map 805 defines a filter, zooming or panning the map 805 always defines a new filter based on the previous filter plus the new domain frame. Each group of icons 810 defines its own unique filter: namely, the filter defined by the current state of the map 805; but with the resulting documents restricted to those associated with at least one of the icons 810 in the group. In this way, a click on an icon 810 can define a filter, since a single icon 810 is simply a group of one.

The general filter display 850 includes filters created for the user. The user-specific filter display 860 includes filters created by the user. The two sets of controls, 850 and 860, can be disjoint or can share controls in the map interface 80.

The general filter display 850 includes general 852 filters, search history 854 filters, and inferred 856 filters. A general 852 filter is a filter predefined by the computing system 20. This includes filters handpicked by human editors to be of general interest to the user population, as well as filters selected algorithmically for having a high frequency of recurrence among the usage patterns of the user population. A search history 854 filter is a filter the current user has applied in the current or previous session possibly without

explicitly instructing the system to remember it. By providing easy access to search history 854 filters, the system allows the user to reapply a filter that he/she created earlier but neglected to add to the user-specific filter display 860.

An inferred 856 filter is a filter selected algorithmically based on the usage patterns of  
5 the current user.

A data-mined 857 filter is a filter created algorithmically by a procedure that analyzes the content and hyperlinks of documents in the page repository 222 to create a set of documents sharing a property. The property may be determined heuristically, e.g. "all 10 documents appearing to relate to cooking recipes." The algorithm to construct such a filter might include the use of Bayesian learning, statistical analysis, and ontologies of words and phrases.

The user-specific filter display 860 is not shown by certain states of the map interface 80. For example, if the computer system 20 cannot determine the correct user profile to apply to the current user, or if a security measure associated with the profile has not been 15 satisfied, the user-specific filter display 860 may be hidden or disabled.

When displayed and active, the user-specific filter display 860 includes filters associated with a user profile. The user can add, modify, or delete these filters, and can assign them to user-defined groups.

Filters that the user can add to the user-specific filter display 860 include: a filter in the general filter display 850; the filter defined by the current state of the map 805; the filter defined by a group of icons 810, which the user can specify by using the pointer symbol; a filter combined from at least two existing filters; and a modified filter which the user chooses to save under a new name.

The modifications that the user can apply to a filter in the user-specific filter display 860 include: renaming the filter; adding, deleting, or reordering elements in its list; and changing the icon class associated with the filter or defining a new icon class for the filter. Properties of the icon class that the user can edit include: its name, its icon face 818, its parent icon class, a textual summary of the document, and any properties displayed in the icon class legend 817.

### ELECTRONIC NOTES

A note document is a document associated with a domain location. It may also be associated with a user profile, or it may exist anonymously. An electronic sticky-note 870 is a representation of a note document displayed on the map 805 in a display location corresponding to a domain location associated with the note document. The note document can contain any form of information that a document in storage 22 can contain. For instance, the note document may contain text, graphics, sound, video, hyperlinks, or a combination thereof. The note document can have its own URL and act as a web page.

The post a note 836 control changes the state of the map interface 80 such that a subsequent click on the map 805 will create a new note document. The note document will be associated with a domain location corresponding to the display location that was clicked, and an electronic sticky-note 870 will appear at said display location and be associated with the domain location represented by that display location.

In one embodiment, having put the map interface 80 in the appropriate state, the user can move document content from outside the client process onto the map 805, thereby initiating a note document creation. The content can be moved by drag-and-drop or copy-and-paste, among other methods appropriate to the computing environment and the media type. For instance, the document content could be a media stream which the computer system 20 begins recording. The content becomes part of the new note document, and the note document is given at least one externally accessible address such as a URL. With the map interface 80 in the appropriate state, the user can create web pages, for example, with one rapid action. In this embodiment, the mechanisms allowing the user to drag-and-drop or copy-and-paste content are provided by the operating system. The terms "drag-and-drop" and "copy-and-paste" are well known in the art.

Several other features of these note documents require description. Users can specify calendar dates and/or times when a document is not to be served to the public, or will expire altogether. When a note document expires, it may be deleted from storage or prevented from appearing in the interface. This allows users to post time-sensitive information at geographic locations. Short lifetime note documents might be used to make an animated icon on the map interface. Such an icon could follow a moving object or a user's approximate path through the domain.

Users can digitally sign note documents to help ensure their authenticity to other users. Public key cryptography, like PGP, is standard in the art and can be used to affect this. The audience of a document can be limited using this same type of public key cryptography or by requiring users to login with a private password that authenticates their identity. The 5 creator of the note document can determine the list of registered users permitted to see a particular note document. Alternatively, the creator could distribute the encryption key needed to open a note document. This allows users to publish note documents to a subscription list.

Users can host their own note documents on private computer systems. Such private 10 computer systems may be licensed copies of part or all of the computer system 20. Such a privately held note document might be protected by security measures. The creator of such a note document can create additional note documents in other instances of the computer system 20, which may be owned by other people or companies. These additional note documents could provide pointers to one or many note documents on the creator's private 15 computer system. These additional note documents might contain a summary of the original note document. Users of one instance of the computer system 20 may have access to certain other instances of the computer system 20. This access is determined by the owner of each instance. This allow many instances of the computer system 20 to participate the hosting and distribution of geographically-located note documents.

Since any media type can be easily put in to a note document, it is easy for the owners 20 of an instance of the computer system 20 to create note documents from data from other computer systems under their control. For example, a store owner can copy their inventory database into note documents in their instance of the computer system 20. This conversion 25 of a store database to geographically-located note documents makes it easy to serve the inventory information to other users interested in the stores physical area.

A user can upload or create a collection of note documents in one action, such as dragging and dropping a folder of documents into the map interface. If the documents contain location information, they can be automatically posted in the map interface. If not, the user can be prompted to select locations for each document.

Such a collection of note documents will be grouped in a filter in the user-specific 30 filter display 860. Examples of such grouped note documents include a collection of

photographs taken on a vacation, a collection of sound recordings taken around a city, a set of data gathered from various sensors, a sequence of events for an newspaper article, or a set of descriptions for a trail guide. A collection might have colored lines connecting the various icons on the map 805, thereby indicating a path that could be followed by a user in the domain.

Such a collection could be created for a user by a service or device. For example, a user's camera might include a GPS or other spatial locating device that imprints each picture with a location stamp. Uploading the pictures is then quite simple: the stamps locate each picture on the map 805. A service might do this on a user's behalf. For example, a hospital might annotate a user's medical record with locations of where the user was treated and post them as a private note document collection for the user and other care providers.

The user can post a note document containing dynamic software such as a discussion board, order entry tools, telephone connect service, or other software-backed tool. A note document posted at the location of a vending machine might have an order entry tool connected to the vending machine that allows users to use a credit card or other payment mechanism to purchase items from the machine. This allows users get physical items without paying cash or even carrying a credit card.

A note document posted at a store might contain a discussion board with text and other media entry tools allowing the general public to engage in a discussion at that location. Such message boards might receive text messaging from portable phones and broadcast them to users viewing the discussion board.

A note document might contain a tool, which, when clicked, causes a user's phone to dial into a service. Such a note document might be posted at a restaurant or theater where telephone reservations are required.

## 25 COMMUNITY FEEDBACK

The map interface 80 can use the community feedback 880 control to show the user information gathered from the behavior of other users. Features of the community feedback 880 control include domain usage feedback 882, word-domain suggestion 884, and word-word suggestion 886.

When the user views a spatial domain, domain usage feedback 882 tells the user how many people have viewed that domain or part of that domain in the recent past. For example, "23 people have viewed this region in the last 18 minutes."

When a user views a spatial domain, word-domain suggestion 884 can tell the user 5 keywords that are relevant to this domain. These words can be gathered by analyzing documents that refer to this region to find the words that occur most in that domain. These words may also be gathered by recording the keywords that other users have entered when viewing this region. The most commonly searched for words can be presented to the user.

When a user enters a keyword query, word-word suggestion 886 can tell the user 10 additional keywords that relate to the keyword(s) just entered. These additional keyword suggestions come from a thesaurus that may be built by recording the sequence of queries entered by other users. If many users enter the same keywords together or in a single session, then those keywords can be considered related. For example, if many users search 15 for "chocolate" and then search for "chocolatier" the computer system 20 can suggest to the next user who enters "chocolate" to try a keyword query for "chocolatier." This suggestion helps users find what they want.

#### *DATA COLLECTION*

The computer system 20 includes a data collection 30 process for gathering new 20 documents. With reference to Fig. 3, the data collection 30 process includes a crawler 36 process, a page queue 34, and a metasearcher 32 process.

#### *CRAWLER AND PAGE QUEUE*

The crawler 36 loads a document over a network, saves it to the page repository 222, and scans it for hyperlinks. By repeatedly following these hyperlinks, much of a networked 25 system of documents can be discovered and saved to the page repository 222. The crawler 36 gathers documents into the computer system 20 in this manner. In one embodiment, these documents are World Wide Web pages available on the Internet. In this case, downloading pages can be done using any of the various Internet protocols, including the HyperText Transfer Protocol (http), the File Transfer Protocol (ftp), gopher, news, wais, and others.

The page queue 34 stores document addresses. The crawler 36, the pioneer 48, and the metasearcher 32 add document addresses. The page queue 34 comprises a database table, the page queue table 340.

The crawler 36 gets document addresses to crawl from the page queue 34. When the crawler 36 loads a previously unknown document, it passes the document to the pioneer 48 process. The pioneer 48 parses the content of the document for hyperlinks to new documents. The pioneer 48 adds any addresses referenced by such hyperlinks to the page queue 34.

The crawler 36 makes use of the fact that the probability of being spatially relevant is correlated with linkage; in other words, pages linked to a spatially relevant page have a greater probability than average of being spatially relevant. Each crawled URL is assigned a spatial relevance. Considering spatial relevance helps the crawler 36 use time and other resources efficiently.

The crawler first crawls pages linked from those pages with spatial relevance greater than a predetermined threshold. After a page has been downloaded and its spatial relevance calculated, its spatial relevance level 342 field can be recalibrated to reflect the actual relevance we found.

#### *METASEARCHER*

The metasearcher 32 initializes the collection of known documents. This initializing step is called "seeding" or "bootstrapping." The computer system may have to be seeded for each domain. For example, separate bootstrapping operations may be used for United States postal addresses and French postal addresses.

The metasearcher queries search engines known to store information appropriate to the domain, such as search engine web sites on the Internet. The human administrators of the metasearcher provide it with a collection of known spatial locations appropriate to the domain. The metasearcher formulates queries based on these spatial locations and directs the queries to the search engines. The results are compared to the collection of known documents and are added if new.

A crawling is complete when all discoverable documents on the network have been found. In practice, this rarely happens over large document collections unless the collections

are extremely static. Thus, since a complete crawling is rarely likely, the speed of the crawl is an important design concern. The speed of crawling is limited by the speed at which new pages are discovered through links on previously downloaded pages. A good way to accelerate this crawling is to query existing search engines that have already crawled at least 5 part of the document collection, which could be the Web. The results given by these search engines are used to bootstrap the data collection 30 process.

In one embodiment, the metasearcher 32 bootstraps its knowledge of the geography of the United States. The process for this bootstrapping comprises six steps. Other domains may require different processes.

10 The steps are a system of levels intended to gather the most useful spatial URLs from existing search engines. Since search engines commonly limit the number of results returned to a single query, searches might not return all the results that we would like to gather. For instance, in a geographic query, this happens with town names like "Boston, MA." In such cases, it is useful to specify other words in the query, such as all the street names in that 15 town.

Major search engines include AltaVista, Fast, Lycos, MetaCrawler, DogPile, NorthernLight. Each engine has a maximum number of results that they will return for a query, even if they have more pages that meet the query. If a metasearch query overflows this number, the metasearcher 32 adds words to the query to squeeze out more URLs.

20 In step 1, the metasearcher 32 queries the search engines with just the town names, e.g. "boston" "cambridge" "new york" "madison" "san antonio".

In step 2, for any town name that resulted in the maximum number of results for that engine, the metasearcher 32 re-queries the search engine with the town and the state, e.g. "boston, ma" "boston mass" "boston massachusetts" "cambridge, ma" etc... "new york, ny" 25 etc ... "madison nj" ... "madison ny" ...

In step 3, the metasearcher 32 switches to a second table, which has more information. The second table includes all the streets in every town in the USA. For any town-state pair that overflows on a particular engine, the metasearcher 32 queries for every street, e.g. "highland somerville" "hancock somerville" "elm somerville" etc.

30 In step 4, the metasearcher 32 adds in state names with the street names, e.g. "highland somerville ma" "hancock somerville ma" "elm somerville ma" etc.

In step 5, the metasearcher 32 adds in street types, e.g. "highland ave somerville" "highland avenue somerville" ... "hancock st somerville" ... "elm st somerville" etc.

In step 6, the metasearcher 32 adds in street types and state names, e.g. "highland ave somerville ma" "highland avenue somerville ma" "highland avenue somerville massachusetts" etc. Few places reach this level.

The page queue table 340 includes a spatial relevance level 342, which helps constrain the crawler 36 to documents that are spatially relevant. When the metasearcher 32 gathers a document, the document is given a level of "0."

## *DATA ANALYSIS*

With reference to Fig 4., the computer system 20 includes a data analysis 40 process for extracting information and meta-information from documents. Data analysis 40 includes a spatial recognizer 42 process, a spatial coder 43 process, a keyword parser 44 process, an indexer 46 process, a spatial document ranking 45 process, and a pioneer 48 process. The role of the pioneer 48 process is described in the section for data collection 30. In the data analysis section, we will repeatedly cite the example of the standard geographic domain for the USA, identified by the standard latitude/longitude but also by postal system addresses, localities, and phone numbers.

## *SPATIAL RECOGNIZER*

As new documents are saved in the page repository 222, the spatial recognizer 42 opens each document and scans the content. It searches for patterns that resemble parts of spatial identifiers. For example, in the standard geographic domain for the USA, patterns include street addresses of the USA postal system, localities, and phone numbers.

In step 422, the spatial recognizer 42 finds candidate spatial data in unstructured text. Candidate spatial data, is called a PSI, for possible spatial identifier.

In step 424, the spatial recognizer 42 parses the text of the candidate spatial data to determine its structure, thereby forming a PSI. We break addresses into a standard set of fields used by the US postal system. Similar formats exist for other postal systems, which would be represented as other domains. The constituent parts of the PSI are identified. Not all may be present in a given document; for localities and phone numbers, only town, state, and possibly ZIP and ZIP+4 are used. The constituent parts include:

House number  
Street prefix (e.g. East, South)  
Street name  
Street suffix (e.g. East, South)  
5 Street type (e.g. Street, Turnpike, Square)  
Town  
State  
Zip  
4-digit zip extension

1 10 PSIs are stored in the spatial lexicon 224 for further analysis. The table for these possible spatial identifiers (PSIs), which in this case is mapped against the standard geographic domain, includes fields for latitude and longitude. Regardless of domain, the table may include fields for spatial coding confidence, number of documents located at this place, status of spatial coding, and sum of relevances of documents located at this place.

15 The relevance scorer 426 assigns a relevance score to the document.

The relevance scorer 426 includes a multiple spatial references partitioner 4262 process. Many documents have multiple spatial references. It might be the case that all the spatial identifiers are relevant to the whole document. An example is a web page listing branch locations of a store chain. However, it can be the case instead that each spatial identifier is only relevant to a proper subset of the page. An example of this is a page giving short reviews of a number of restaurants. Such a page is a multi-part document.

20 Multi-part documents present a problem when searching the document collection by keyword. Were the document to be keyword indexed as a whole, a word in one part of the document would be indexed as though it were relevant to addresses in a different part of the document, when in fact the word may not be relevant to that part.

25 To detect multi-part documents, the multiple spatial references partitioner 4262 invokes the multi-part cluster measurement 42625 process. The multi-part cluster measurement 42625 process first rejects any document with fewer than some number of addresses (usually 5) or which is shorter than some number of words (perhaps 200). The multi-part cluster measurement 42625 process computes an array containing the fractional positions of each PSI in the page. For instance, an address that begins at the 200th word in a 1000-word document is at fractional position 0.2. We then apply a clustering statistic such as the Gini coefficient to produce a clustering score that expresses how concentrated the

addresses are on the page. Documents with low clustering score (indicating that the addresses are evenly spread out) are likely to be multi-part documents. The threshold for the maximum clustering score is determined empirically and may vary for each domain.

The multiple spatial references partitioner 4262 partitions the document into segments that contain one PSI each, using the PSIs as boundaries, as follows. The nth segment, containing PSI n, begins at the word following the end of PSI n-1, and ends at the word before PSI n+1. For n=1, the segment begins at the first word. For the last PSI on the page, it ends at the end of the page.

Each segment then has the title portion of the document added to it. The tag recognizer 442 provides one way of determining the title portion of a document.

The segment is stored in the page repository 222 to be separately indexed. The unsegmented page is retained, so that when a segment is found as a search result, the full document can be returned, with an anchor placed at the beginning of the segment so that the document can be scrolled to the segment before presenting it to the user.

#### SPATIAL CODER

To further analyze the PSIs, the spatial coder 43 process runs several processes that associate domain locations with various identifiers in the document content. In the standard geographic domain, we can associate latitude/longitude points or bounding polygons with identifiers; this process is known as geocoding. If no latitude/longitude can be matched to a PSI, the spatial coder 43 marks it misrecognized. Otherwise, the spatial coder 43 turns the PSI into a known spatial identifier, or KSI. This completes the entry in the spatial lexicon 224 described above.

The spatial coder 43 for the standard geographic domain for the USA includes an address encoder 432, a locality encoder 434, and a phone number encoder 436.

With reference again to the standard geographic domain for the USA, addresses are considered the best match. Thus, if a page has addresses in it, simple place names like "Cambridge, MA" and phone numbers are not used to spatially code the page. A page can have multiple KSIs, but that reduces its spatial relevance (see spatial document ranking 45), so we look primarily for pages with only a few highly focused KSIs. A focused KSI means that the spatial coder 43 associates a small area in "lat/long space" (space identified by

latitude and longitude) with high certainty. Thus, for example, a phone number associates with a region the size of a telephone exchange, which is at least several square miles, but a postal address associates with a "rooftop" sized region usually represented by a point in the middle of the hypothetical rooftop. If a phone number and an address in a document both agree on the location of the page, we can improve the ranking of the document (spatial document ranking 45).

Address encoder 432: Postal addresses in the USA and other countries can be associated with small geographic regions, usually the size of a building. Standard geocoding procedures approximate this by a point. Given a PSI like this, for instance:

1 10            77 massachusetts ave|cambridge|ma|02139  
the associated lat/long can be discovered by feeding the text string into any standard address geocoding product. Examples include Etak's Eaglecoder, Sagent's GeoStan, and ESRI's ArcINFO geocoding plug-in. The output of Etak's Eaglecoder looks like this:

1 15            <command line interface> jrf@raag:~\$ mc/lib/etak/rie -b  
<input text of PSI> 77 massachusetts ave|cambridge|ma|02139  
<output of geocoder> 77 MASSACHUSETTS AVE,CAMBRIDGE,MA,02139,42.358968,-  
071.093997

The third line of the output contains lat/long information to associate with this address. Thus, this PSI can be converted into a KSI.

2 20            Locality encoder 434: Place names, like "Boston, MA" and "Washington Monument," are listed by the US Census along with the latitude longitude of the center of the place. This makes it easy to geocode them. The locality encoder 434, similar to the address encoder 432, searches for candidate strings that could be town and state names. The locality encoder 434 differs, however, in that it looks up the town name in a database of all known towns in the  
2 25            United States 2262, and rejects the town name if it does not appear.

Phone number encoder 436: The phone number encoder 436 converts phone numbers to geographic locations by looking up the area code and exchange in a phone-to-place table 2266. The phone-to-place table 2266 maps area code-exchange pairs to town name-state name pairs. This pair is then treated as a locality name, except that its relevance score is lowered by a small constant number (determined heuristically) to reflect the fact that towns obtained in this way are somewhat less valuable than towns that have been mentioned by name. A single telephone company central office may cover multiple towns, especially in

suburban locations; there is a chance that the phone number is actually located in a neighboring town.

#### *SPATIAL MEANING INFERENCE*

The spatial coder 43 includes a spatial meaning inference 438 process, or SMI 438, which can perform a special type of spatial coding. The SMI 438 can deduce a spatial relevance for terms (words and phrases) based not on a semantic interpretation but on statistical properties of appropriate portions of the spatial-keyword document index 505.

Certain words and phrases correspond to geographic locations but are not recorded by any existing geocoding services. To discover these geographic relations, the SMI 438 statistically analyzes the correlation of candidate words and phrases with KSIs. The SMI 438 uses the premise that if a phrase occurs mostly in documents with addresses in the same place, then the phrase is probably also about that place. For example, "the big apple" occurs on many pages with the words "New York, NY" and addresses in New York City. The SMI 438 can deduce that "the big apple" is also about New York City.

The SMI 438 deduces spatial relevance as follows. The spatial-keyword document index 505 contains a tree for each indexed term, i.e. each term in word lexicon 225. For each word in a given string, the SMI 438 examines the tree associated with that word. The examination includes invoking the imbalance measurer 439 to measure a degree of imbalance in the structure of the tree, which, since it is a trimmed version of the spatial document index 503, may have significant imbalance as a result of trimming. The imbalance measurer 439 is described below. Broadly speaking, and as will be described in more detail, if enough terms in the string have trees which have similar imbalances, the SMI 438 associates the string with the spatial regions described by the imbalanced portions of said trees.

Returning to an earlier example, each word in the phrase "the big apple" appears in many documents. Performing a search over a spatial-keyword document index 505 for that phrase without specifying a bounding box will find a large "peak" in the number of documents near New York City. This is evidenced by the degree of imbalance in the trimmed result tree. The tree resulting from the intersection of these three words has many branches in the latitude-longitude region covering New York City. This tells us that pages with these three words next to each other are probably referring to this lat/long region.

We call such words and phrases "geographic phenomena."

A tree address is defined as follows. Given a spatial-keyword document index 505, any node or leaf in the index trees can be identified by a set of values indicating the sequence of child node numbers that must be traversed to reach that node. For example, in a binary tree, the tree address 0110 specifies the node found by starting at the root node and going to the first child's second child's second child's first child. In a 16-way tree, the tree address written in hexadecimal as "0x4f8" specifies the node found by starting at the root node and going to the fifth child's sixteenth child's ninth child.

To measure the "peakiness" of a particular phrase without using a spatial-keyword document index 505, the imbalance measurer 439 first computes a "standard peakiness" of average words and then compares candidates to that. In one embodiment, the imbalance measurer 439 computes the standard peakiness by picking a random sampling of words and, for each of those words, computes the 2-dimensional variance of the points referred to by documents that contain the word. Documents that are particularly relevant to a word can be given extra weight in computing the variance, e.g. a highly relevant document can be scaled linearly so that it appears to represent multiple documents at that location. Given this random set of variances, the imbalance measurer 439 computes the average variance. The average variance can be used as a baseline to detect a geographically relevant phrase or word. Any word or phrase with a variance much smaller than the baseline is a geographic phenomenon.

Use of the spatial-keyword document index 505 simplifies the SMI 438 dramatically. Since the trees in the spatial-keyword document index 505 already span all the documents known to the computer system 20, the SMI 438 can detect a geographic phenomenon simply by considering the set of tree addresses of leaves in a trimmed result tree. For example, given a candidate word or phrase, the SMI 438 queries the spatial-keyword document index 505 to get the trimmed result tree for this word or phrase and performs the following operation on this list of addresses.

From the tree, the SMI 438 creates a list of the tree addresses of every leaf. Starting at the first digit in all the addresses, the SMI 438 finds the most common branch number at this level (i.e., for this digit). The branch indexed by this digit is called a "candidate fork"

because it is a fork of the tree, pointing in the direction of the candidate location. The SMI 438 computes the fraction of the addresses that follow the candidate fork at that level.

At the next level, the SMI 438 considers all addresses that took the candidate fork in the last level and once again finds the most common fork direction, using it as the next fork direction. The SMI 438 again computes the fraction of addresses still following the candidate fork.

The SMI 438 repeats this until the percentage of addresses still following the candidate fork falls below a predetermined threshold adjustable by the operators of the computer system 20. The particular threshold may be adjusted for each domain. Adjusting the threshold adjusts the quality of matches that are considered. It is set empirically.

For example, for simplicity of explanation consider a binary tree whose nodes divide a domain space into rectangles, and consider these four addresses that fork together for several levels:

101111010101111  
1011101011101010  
1011101011101111  
1011101011101101

15 Level 1: forked 1 = 100%

20 Level 2: forked 0 = 100%

Level 3: forked 1 = 100%

Level 4: forked 1 = 100%

Level 5: forked 1 = 100%

Level 6: forked 0 = 75%

25 Level 7: forked 1 = 75%

Level 8: forked 0 = 75%

Level 9: forked 1 = 75%

Level 10: forked 1 = 75%

Level 11: forked 1 = 75%

30 Level 12: forked 0 = 75%

Level 13: forked 1 = 75%

Level 14: forked 1 = 50%

Level 15: forked 0 = 25% -- below 50% threshold.

These tree addresses suggest that the word is 100% relevant to a region defined by the rectangle in the spatial index tree 502 by the address 10111, and 75% relevant to the rectangle 10111010111.

If a particular word is rare, i.e. occurs only a few times in the entire page repository 222, but its appearances are highly correlated with geographic identifiers in the same place, then that word might be associable to a point location. For example, the word "EVOO" is the name of a restaurant in Somerville, MA, USA. The word "EVOO" appears only a few times in the entire corpus. Most of these times it appears on a page with the address for the restaurant. The other times, it appears on pages reviewing the restaurant. Given the strong correlation of "EVOO" with the restaurant's address, we can geocode the word "EVOO" with the same latitude/longitude point. This enables us to geocode the other pages with that same point. The latitude/longitude point is transmitted from one page to the other pages through the word link "EVOO."

Note that the spatial meaning inference 438 process is not usually able to associate a phrase with a location as focused as a point. Bounding polygons are a more common result. The main purpose of geocoding these phrases is to improve the ranking of documents, discussed in the section on spatial document ranking 45.

#### **KEYWORD PARSER**

Non-geographic search terms (keywords) are identified as follows. As the documents are saved to the page repository 222, a keyword parser 44 process opens each document and scans its keywords. These keywords are stored in a database table called word\_instances 227, which includes the fields: wordID 2272, docID 2274, and word-doc relevance float 2276. The word\_instances 227 table associates a given keyword with a set of documents containing it.

The WordID is a number that replaces the string of characters in the word. This reduces storage requirements and allows us to treat a phrase like "the big apple" as a single database entry. The word lexicon 225 is a database table that acts as the dictionary of all words and their corresponding WordIDs. The word lexicon 225 table includes the fields: word 22621; wordid 22623; and word\_occurrences 22625.

The keyword parser 44 includes a tag recognizer 442 for parsing documents that contain tagged text such as SGML or the related standards HTML and XML. Tag recognizers for various document standards are well known in the computing art and can even be a feature of the operating system.

5 Methods standard in the art may be used to index a document for phrase searching, this allows a user to issue a query for a set of words close together or immediately adjacent in documents.

#### *SPATIAL DOCUMENT RANKING*

Given the potentially vast amount of information, document ranking is very important. Results relevant to the user's query must not be overwhelmed by irrelevant results, or the system will be useless.

10 The spatial document ranking 45 process produces a ranking of documents that includes evaluations of document-to-place relevance 452, document-to-word relevance 454, and abstract quality 456. Evaluations are combined into a floating point number indicating 15 the relevance of each document to the query.

15 The document-to-place relevance 452 score indicates a document's relevance to a domain location, where the domain location is described by a PSI or KSI within the document. The following is a method of considering the relevance of one SI (spatial identifier, which might be a PSI or a KSI) to one document. It is possible to compute this for 20 several different SIs in the same document. These SIs can be combined if they all refer to the same geographic region. For example, a document might have an address and a phone number that we can geocode. If the address is to a point nested inside the phone number's area, then we can improve the geographic relevance of the document to that address. The boost in relevance might be affected by handcrafted weights chosen for the different 25 circumstances in which multiple SIs can combine on a page. This improvement is secondary to the relevance computed by the following method.

#### *DOCUMENT-TO-PLACE RELEVANCE*

30 The document-to-place relevance 452 score includes the following scores: position in page 4521, distance from end 4523, number of other SIs 4525, in sentence 4527, and emphasis 4529. (See Appendix A)

The position in page 4521 score is a heuristic function, calibrated from large numbers of observations of SIs. It assigns a score on the premise that SIs appearing earlier in a document are likely to be more relevant. Distance may be measured in characters or bytes. SIs that appear "above the fold" (on screen when a page is first loaded, without having to scroll) are considered most relevant.

5 The distance from end 4523 score gives the document-to-place relevance 452 score a slight boost if the SI occurs at the footer of the document; this partially counteracts the low score assigned to it by the position heuristic.

10 The number of other SIs 4525 score is a heuristic function that dilutes the relevance of a SI based on how many other SIs are in the same document. Documents with large numbers of addresses tend to be lists, where any individual address has a low probability of being relevant to the document.

15 The in sentence 4527 score gives a slight boost to SIs that are free-standing, as opposed to being mentioned in a sentence.

20 The emphasis 4529 score reflects the degree of emphasis of the SI text, including being in boldface, large type, or in the page's title. This score takes the form of a decimal number where 1.0 is assumed to be standard (neither de-emphasized or emphasized); lower numbers indicate lack of emphasis (such as small text) and higher numbers indicate prominence.

#### 25 DOCUMENT-TO-WORD RELEVANCE

20 The document-to-word relevance 454 score indicates the relevance of a particular word to a particular document that contains it. Means for measuring the relevance of a word to a document are well known in the art. For instance, see S.E. Robertson and K. Sparck Jones, "Simple, proven approaches to text retrieval," University of Cambridge Computer Laboratory technical report, May 1997.

Phrase searching may also affect document relevance. This type of relevance is typically computed on-the-fly at the time of a user's query for a particular phrase. There are methods standard in the art for computing this type of relevance.

### *ABSTRACT QUALITY*

The abstract quality 456 score represents document value independent of a given word or place. There are several ways to measure this, including the number of pages that link to the document, the number of times people click on the document when it is served as a search result, and the number of other documents that refer to the same words and places – that is, if it is a document like many others, its abstract value might be considered low, independent of the particular words it contains.

The abstract quality 456 score include components for network connectedness 4562 and a manual updates 4564. Network connectedness 4562 is computed from the probability that the page will be chosen by a random crawl of the web. This probability is then mapped to a score. The particular mapping chosen depends on the size of the document collection in the page repository 222, since the probability of finding any given document is inversely proportional to the collection size.

The manual updates 4564 score is designed to incorporate the input of human editors. The editors can craft rules that adjust the abstract quality 456 of particular documents. For example, they can weight all documents within a particular site as better than other documents simply by increasing their document quality measures. They might do this with a site that itself is a careful product of human editors, such as Zagat.com.

The abstract quality 456 score is stored in an abstract\_document\_quality 228 table, which includes the fields doc\_id 2281 and document\_quality 2283. The doc\_id 2281 field is a foreign key referencing the doc\_id 2221 field in the page repository 222.

### *INDEXER*

The indexer 46 analyzes documents to prepare data structures that accelerate the search 50 process. The indexer 46 includes a spatial indexer 462, spatial-keyword indexer 465, and a tree degree converter 466.

#### *SPATIAL INDEXER*

With reference to Fig. 7, the spatial indexer 462 creates a spatial index 502 and a spatial document index 503 for a domain space. The spatial index 502 is a binary tree. The spatial document index 503 is a tree that is based on the spatial index 502, but may be of a higher degree than 2 (the degree of all binary trees).

The spatial indexer 462 in step 4621 gathers a collection of all domain locations referenced by a document in the page repository 222, then creates a root node for the spatial index 502 tree in step 4622. The spatial indexer 462 passes the root node and the collection to step 4624, which marks the beginning of the recursive spatial indexing subroutine (or RSIS) 4620.

5

In step 4624, the RSIS 4620 receives a node and a collection. The RSIS 4620 examines the collection in step 4625 to determine whether the collection contains more than one element. If it does not, the RSIS 4620 associates the current node with the one element's domain location in step 46295 and goes to step 4629, returning control to the routine that invoked it. Otherwise, the RSIS 4620 proceeds to step 4626, where the RSIS 4620 spatially divides the collection along spatial divider D into collections L and R, such that L and R are as equal in number as possible. If the domain space is a plane, the spatial divider D is a line in the plane. If the domain space is in three dimensions, the spatial divider D is a plane through 3-space. In general, if the domain space is of X dimensions, the spatial division is a manifold of dimension X minus one. The RSIS 4620 in step 4626 also stores the criteria for the spatial divider D in node N. Thus, each node contains criteria that divide a master collection of locations into two sub-collections.

10

15

20

The RSIS 4620 in step 4626 also creates a left node and a right node on the node passed to step 4624. This creates a fork in the binary tree that will act as an index. The tree as a whole becomes the spatial index 502.

25

The RSIS 4620 becomes recursive by invoking itself on each of the sub-collections. Specifically, in step 4627 the RSIS 4620 passes sub-collection L and the current left node to step 4624, while in step 4628 the RSIS 4620 passes sub-collection R and the current right node to step 4624. The RSIS 4620 repeats until every collection has been divided into collections of single elements, which are associated with childless nodes. All other nodes have division criteria and two nodes descending from them.

30

After the spatial indexer 462 builds the spatial index 502 tree, which indexes the points referred to in a corpus of documents, the spatial indexer 462 builds the spatial document index 503 by extending a copy of the spatial index 502 tree to cover multiple documents that refer to the same spatial point. The spatial indexer 462 invokes a tree degree

converter 466 to make a version of the spatial index 502 that is represented in a tree of degree k.

The extension of the spatial index 502 produces new branches that no longer reflect spatial divisions but instead reflect partitions of the documents referring to that point. In particular, instead of the nodes including criteria that define spatial divisions within the domain (as the nodes inherited from the spatial index 502 continue to do), the nodes added after the extension include criteria for branching within the space of the docID 2221 numbers of the documents. Partitioning based on a key value (such as the docID 2221) of a database table is standard in the art. Such a partitioning produces a k-way tree on the documents using their docID 2221 numbers as a key.

#### *DEGREE K*

An important optimization of an index tree of degree k involves the selection of k. A k-way branching structure must be chosen before building or storing the trees. K could be as low as two and as high as a few thousand or tens of thousands, depending on the number of documents and possibly the underlying computing platform. A tree of degree k can index  $(k)^L$  documents in L levels.

A large value for k makes it faster and more storage efficient to deal with keywords that appear in only a few documents. If the number of rare words in the page repository 222 is large, a large value for k is more storage efficient than a smaller one. However, a smaller value of k can be more search efficient, as it allows a traversing process (in response to a query) to ignore branches of the tree that fail its constraints.

The selection of k is an empirical process that may be performed for every set of documents, in the page repository 222, to be indexed. It is influenced by hardware limits, such as the number of bits handled by a single processor instruction and the number of blocks loaded by the disk drive. The most important factor in choosing k is the word-frequency distribution. Keyword lexicons for web pages, for instance, show a huge number of words that appear in only one or two documents, but more common words appear in many documents. These common words produce a "fat tailed" distribution. The exact shape of the distribution for a particular set of documents determines the optimal k. Given a value for k,

it is a simple calculation to count the number of bytes used to store the word trees for a particular lexicon and set of documents.

#### *TREE DEGREE CONVERTER*

The tree degree converter 466 is a function that accepts parameters including a binary tree and an integer k, and returns as its output a tree of degree k incorporating the structure and data of the binary tree. Methods for this conversion are known in the computing art.

#### *SPATIAL-KEYWORD INDEXER*

The spatial-keyword indexer 465 builds a spatial-keyword document index 505 responsive to queries for documents. The queries can have keyword criteria, spatial criteria, or both.

The spatial-keyword indexer 465 gathers all domain locations referenced by documents in the page repository 222.

The spatial-keyword indexer 465 uses the spatial document index 503 generated by the spatial indexer 462. The spatial document index 503 is a k-way tree on this list of documents. The spatial-keyword indexer 465 copies the spatial document index 503 to create a keyword tree 506 for every keyword. For each keyword tree 506, the spatial-keyword indexer 465 trims away all documents that do not contain that particular keyword. If, after the document trimming, the subtree depending from a node of the keyword tree 506 does not contain a document, the spatial-keyword indexer 465 removes that node (and therefore its subtree).

The spatial-keyword indexer 465 creates for each keyword a minimal keyword tree 506 that relates the keyword to the corpus of documents in the page repository 222. Furthermore, the spatial-keyword indexer 465 ensures that one branching structure is common to all keyword trees as well as to the spatial document index 503 tree.

25

#### *SEARCH*

With reference to Fig. 5, the search 50 process responds to queries with a set of documents ranked by relevance.

A lexical tree 508 is any copy of the spatial document index 503 tree, possibly trimmed. Thus, every keyword tree 506 is a lexical tree 508, as is the spatial document index 503 tree itself. Also, any filter can be expressed as a lexical tree 508, since a filter determines a set of documents, and any set of documents determines a trimming of the spatial document index 503 tree. Thus, lexical trees 508 can be built to index arbitrarily complex sets of documents.

The search 50 process uses the spatial document index 503 and spatial-keyword document index 505 to find documents that refer to a given set of domain locations or regions, and documents related to a given set of keywords existing in the word lexicon 225. The search 50 process can also find documents using a lexical tree 508, such as might represent a filter. Thus, the search 50 process can respond to queries that seek documents according to spatial domain criteria, keyword criteria, filters, or any combination thereof. Furthermore, the search 50 process can invoke the document ranker 56 process to rank the result set of documents by relevance to the query terms.

The search 50 process answers queries via the procedure in Fig. 6. A query includes at least one of the following: a bounding region specifying a closed shape (typically a polygon in two dimensions), words, phrases, and layers. The bounding region can be the domain frame from the map interface 80.

For each element in the query, the search 50 process loads a copy of the appropriate tree, determined as follows. If a bounding region is specified, step 703 loads the spatial document index 503. If keywords are specified, step 702 load the spatial-keyword document index 505 tree for each keyword. If a phrase is specified and the phrase is not a single entry in the word lexicon 225, then step 702 loads each word's spatial-keyword document index 505. If a phrase is specified and is a single entry in the word lexicon 225, then step 702 need only load that phrase's spatial-keyword document index 505. If a layer is specified, its name identifies the appropriate lexical tree 508, which is loaded by step 702.

The search 50 process counts the number of leafs of each of these trees. At step 703, the search 50 process estimates the approximate number of leafs in the spatial document index 503 bounded by the query bounding region, by multiplying the area of the bounding

region by the average density of points in the corpus. At step 704, these numbers are used to order the trees in a list, with the smallest tree first.

At step 705, this smallest tree is re-labeled as the result tree and will be trimmed to generate the final result tree. For each node that exists in the partially trimmed result tree, the search 50 process checks all the trees to see if they also contain that node. In steps 708 and 712, the search 50 process checks the trees in list order. If any tree lacks that node, the search 50 process stops checking and in step 709 deletes the subtree below that node in the result tree. (See Appendix B) Steps 710 and 711 traverse the tree. The search 50 process continues checking all the nodes in the result tree until only leaf nodes remain. These leaf nodes represent the result set of documents. Step 713 returns the result tree.

The leafs of spatial-keyword document index 503 trees have word relevances and lists of the positions and contextual emphasis of the words in each document. The spatial document index 503 has spatial relevances for each document. The lexical tree 508 for each layer may have an abstract document quality 456 for some documents. These relevances are combined for each document in the result set. The combination procedure might be averaging, summing, or a weighted average.

A second process might compute adjustments to the document relevances by considering the emphasis and proximity of multiple query words within the documents. This standard procedure simply gives higher relevance to documents in which the query words appear closer together.

The final result list of documents might be sorted for return to the user. The sorting procedure might extract only a portion of the documents with the highest relevance.

#### *DOCUMENT RANKER*

Document ranker 56 combines various relevance scores for each document in a result set and sorts the documents by this combined relevance. The combination function may be an averaging or a weighted sum or some other combining function tailored to the various relevance scores used. The document ranker 56 may take streams of sorted result sets from several database systems and merge sort them to produce a new result set.

### *ICON RANKER*

The icon ranker 57 receives a sorted list of results from the document ranker 56. To present this list to the user who requested the documents, the icon ranker 57 aggregates overlapping icons according to the manner described in the section on Icons. This list of 5 aggregated icons is presented to the user with sublists next to each icon 810. These sublists identify the documents aggregated into that icon 810.

The icon ranker 57 groups documents into icons 810 as follows. The icon ranker 57 takes the first document from the sorted result list and makes it the first icon 810 in the icon list. For each subsequent document having a tentative display location in the result list, the icon ranker 57 examines whether an icon 810 situated at the tentative display location would collide with any icon 810 already in the icon list. If a collision occurs, the icon ranker 57 associates the colliding document with the existing icon. If no collision occurs, the icon ranker 57 adds an icon 810 to the icon list and associates the current document with said icon 810. This procedure may terminate whenever the number of icons reaches the lesser of a maximum number determined by the user or a predetermined number that is a customizable 10 operating parameter of the computer system 20.

If a document is topically affiliated with a particular icon class, the icon ranker 57 assigns an icon face 818 from said icon class to the icon 810 that will represent the document. If multiple icon classes are affiliated with documents represented by a single icon 810, the icon ranker 57 may select one of the said icon classes to assign to said icon 810 or 15 may assign a new icon class built to reflect said multiple icon classes.

### *USER PROFILES*

The user profiles 65 process manages information specific to user accounts. The information may include descriptions of how users have interacted with the computer system 20 in the past. Other elements that might be recorded include default location to display to the user when beginning an interaction, set of previously collected layers, set of previously posted note documents, previous searches, and previous click patterns or behavior. Part or all 25 of this information may be made directly viewable and editable by the user.

The user profiles 65 process also allows a user to log into the computer system 20 with a user name and possibly a password. The user name identifies the user with a user 30

account, as is common in the art. The map interface 80 can include account login entry controls 861, including a prompt for account login 862, a data entry control 863, and a submission control 864.

#### *DATA PRESENTATION*

5       The data presentation 60 process manages the state of the map interface 80 for each user session. As the user changes the state of the map interface 80 – for instance, by issuing queries, selecting controls, and generally utilizing the interface tools – the data presentation 60 system keeps track of these changes and their sequence. This recorded history enables querying within previous result sets. For example, a user can query for documents referring 10 to "shoes" in "cambridge, ma," and in a subsequent interaction, the user can filter this set of documents further by requesting only those documents that contain the word "store." This results in a list of documents containing "shoes" and "store" and referring to "cambridge, ma." The user could then zoom out to see a larger region with these document still displayed 15 in the map. To see new documents that might fit the keyword query in this larger domain, the user can re-issue the query.

Similarly, the user could combine a set of documents with another set of documents selected by a different query.

Any number of subsequent filter operations or result set combinations can be performed, limited only by the storage resources of the computer system 20 or, optionally, by 20 parameters built into the computer system 20, as for performance reasons. The data presentation 60 system keeps track of filter operations by a given user so that the computer system 20 can present the correct set of documents to the user at each subsequent interaction.

#### *SERVICE COLLECTOR*

The service collector 24 includes a proxy through which the user interface server 62 and the portal server 70 communicate with the processes for data presentation 60, search 50, 25 and user profiles 65.

#### *PORTAL SERVER*

The computer system 20 includes a portal server 70 process. The portal server 70 offers at least some of the services of the computer system 20 through remote procedure calls

and other network protocols. This allows the services, data, and tools of the computer system 20 to be delivered through public portal systems or directly to individuals. Examples of companies offering public portal systems include Yahoo! Inc. of 3420 Central Expressway, Santa Clara, CA, and Sprint PCS of PO Box 8077, London, KY.

5      *ALTERNATE EMBODIMENTS*

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

1 0

**APPENDIX B**

to test the existence of a node address N in a tree T:

```
if (T is a keyword/layer tree) {  
    nodeQ = pointer to root node of T;  
    foreach $step in node address N {  
        next nodeQ = pointer to child number $step from previous nodeQ;  
        if ( nodeQ is a valid child ) {  
            continue foreach loop;  
        } else {  
            exit with return value = "false";  
        }  
    }  
    return "true"; # the loop exited without hitting a nonexistent child  
}
```

```
if (T is the spatial tree) {  
    in the following, polygonP is the bounding region given by the user;  
    nodeQ = pointer to root node of T;  
    foreach $step in node address N {  
        next nodeQ = pointer to child number $step from previous nodeQ;  
        if ( region below nodeQ overlaps polygonP) {  
        } else {  
            exit with return value = "false";  
        }  
    }  
    return "true"; # the loop exited without hitting a division outside  
    # the query's bounding region  
}
```

**APPENDIX A**

psuedocode excerpt for assigning relevance of document to place

# Heuristically-determined parameters:

5

# \$emphasis\_bonus\_modifier determines the importance of the emphasis bit.

# \$sentence\_penalty\_modifier determines the importance of the in\_sentence

# bit

10

# \$sp\_full\_point: the position after which the sentence penalty fully

# applies

# \$sp\_transition\_point: the position after which the sentence penalty

# starts to apply; it goes from 0 at this position to

# \$sentence\_penalty\_modifier at \$sp\_full\_point

# \$end\_bonus\_size: maximum number of characters from the end of the

# document at which the end-of-document bonus applies

20

# \$end\_bonus\_max: the maximum relevance value for which the end-of-

# document bonus applies

# \$end\_bonus\_multiplier determines the weight of the end-of-document bonus

25

# Start with the position heuristic function. This is a nonincreasing

# function which is normalized to 1 for position 0. It decreases slowly up

# to some position p\_f which is the average position of the "fold", that

# is, the place where the end of the visible area of a typical document

# occurs when it is first displayed to a user. For positions

30

# greater than p\_f it decreases more quickly, but levels off for large  
# positions. The exact form is determined heuristically by manually  
# assigning a score to a large number of instances of PSIs in typical  
# documents and fitting a function to these scores.

5

\$relevance = &position\_function(\$pos);

1 10 # Bonus for being bold, large font, in title, etc. \$emphasis is a  
# heuristic function of the PSI which was assigned on how emphasized it  
# is.

\$emphasis\_bonus = \$emphasis\_bonus\_modifier \* \$emphasis;

15 # Penalty for being in a sentence, e.g. "We would like to announce the  
# availability of several of our products through the Hopkinton Drug  
# Store, 52 Main Street, Hopkinton, MA 01748."  
# No penalty is assigned for PSIs in the first \$sp\_transition\_point  
# characters, going up to the full penalty after \$sp\_full\_point  
# characters.

20  
if (\$pos > \$sp\_full\_point) {  
\$sentence\_penalty = \$sentence\_penalty\_modifier \* in\_sentence;  
} else {  
25 if (\$pos > \$sentence\_penalty\_transition\_point) {  
\$sentence\_penalty = \$in\_sentence \* \$sentence\_penalty\_modifier \*  
(((\$pos-\$sp\_transition\_point)/  
\$sp\_full\_point-\$sp\_transition\_point));  
} else {  
30 \$sentence\_penalty = 0.0;  
}  
}

{

\$relevance += \$emphasis\_bonus - \$sentence\_penalty;

5

# Bonus for being at end of document for long documents. It is

# proportional to

# how low the relevance already is, so that already highly

# scoring PSIs don't receive a bonus for being at the end.

10

# This is before the number of PSIs function so that it will be

# depressed by that function (and the last PSI in a big list won't

# score too high.)

15 if (\$size - \$pos < \$end\_bonus\_size && \$relevance < \$end\_bonus\_max) {

\$relevance += (\$end\_bonus\_max - \$relevance) \* \$end\_bonus\_multiplier;

}

20 # Now depress the above score based on how many other PSIs

# appear on the page.

# num\_psi\_function(\$num) is a function which determines how much less

# valuable a PSI is when it occurs together with other PSIs:

# It is nonincreasing, and is one for \$num = 1; it decreases

# quickly for small \$num, and more slowly for large \$num.

25

# This function is determined heuristically as described above for the

# position function.

\$relevance \*= &num\_psi\_function(\$num);

1      **WHAT IS CLAIMED IS:**

- 2            1. An interface program stored on a computer-readable medium for causing a  
3            computer system with a display device to perform the functions of:  
4                accepting search criteria from a user including a free text entry query and a domain  
5                identifier identifying a domain;  
6                in response to accepting said search criteria from the user, retrieving a plurality of  
7                record identifiers each of which identifies a corresponding record which: (1) has associated  
8                therewith a location identifier that locates it at a specific location within the domain  
9                identified by the domain identifier; and (2) contains information that is responsive to the free  
10          text entry query;  
11              displaying a representation of said domain on the display device; and  
12              displaying on the display device a plurality of icons as representations of the records  
13              identified by said plurality of record identifiers, wherein for each of said plurality of record  
14              identifiers, a corresponding one of the plurality of icons is displayed within said  
15              representation of the domain that is being displayed on the display device, the corresponding  
16              icon for each of said plurality of record identifiers being positioned within the representation  
17              of the domain at a coordinate within the domain that corresponds to the location identifier for  
18              the corresponding record.
- 1            2. The interface program of claim 1 wherein the domain is a geographical region and  
2            said representation is a multi-dimensional map of the geographical region.
- 1            3. The interface program of claim 1 wherein said representation is a two-dimensional  
2            map of the geographical region.
- 1            4. The interface program of claim 2 wherein accepting input further comprises  
2            accepting a designation by the user of a designated category and wherein each of the records  
3            corresponding to the plurality of retrieved record identifiers also includes information that  
4            falls within the designated category.

1        5. The interface program of claim 4 wherein accepting said designation by the user of  
2        a category comprises presenting to the user a list of predefined categories and accepting as  
3        the designated category a selection by the user from that list.

1        6. The interface program of claim 3 for causing the computer to perform the further  
2        functions of:

3              after displaying the corresponding icon for each of the plurality of record identifiers,  
4        accepting further search criteria from the user, said further search criteria selected from the  
5        group of search criteria types consisting of a domain identifier input type, a free text entry  
6        query input type, and a category type;

7              in response to accepting said further search criteria from the user, retrieving a subset  
8        of said plurality of record identifiers, wherein said subset of said plurality of record  
9        identifiers identifies all record identifiers among said plurality of record identifiers that fall  
10      within said further search criteria;

11         displaying a two-dimensional map of a revised geographical region on the display  
12      device that is responsive to said further search criteria; and

13         for each of the record identifiers of said subset of plurality of record identifiers,  
14      displaying a corresponding icon within said displayed map, the corresponding icon for each  
15      of the record identifiers of said subset of said plurality of record identifiers being positioned  
16      within the displayed map at a coordinate that corresponds to the location identifier for the  
17      corresponding record.

1        7. The interface program of claim 6 for causing the computer to perform the further  
2        functions of storing as a filter the first-mentioned search criteria in combination with said  
3        further search criteria, wherein said stored filter is retrievable for use by the user in  
4        specifying a future search through the interface.

1        8. The interface program of claim 7 wherein the first-mentioned search criteria in  
2        combination with said further search criteria is an ordered sequence of inputs and wherein the  
3        stored filter is the ordered sequence of inputs preserving the order of the sequence of inputs.

1        9. The interface program of claim 3 for causing the computer to perform the further  
2        functions of:

3              presenting to the user via the display device a map; and  
4              enabling the user to input said domain identifier as part of the search criteria by  
5              interacting with the displayed map.

1        10. The interface program of claim 3 wherein said plurality of icons include an icon  
2        of a first icon class and an icon of a second icon class, wherein the icon of the first icon class  
3        has first visual characteristics and the icon of the second icon class has second visual  
4        characteristics that are different from the visual characteristics associated with the first icon  
5        class.

1        11. The interface program of claim 10 wherein at least some of the records identified  
2        by said plurality of record identifiers are of a first type and at least some of the other records  
3        identified by said plurality of record identifiers are of a second type and wherein records of  
4        the first type are displayed using the icon of the first icon class and records of the second type  
5        are displayed using the icon of the second icon class.

1        12. The interface program of claim 3 wherein at least one of the icons of the plurality  
2        of icons represents multiple of the records identified by said plurality of record identifiers,  
3        each of said multiple of the records having a location identifier that locates that record within  
4        a neighborhood about a central location.

1        13. The interface program of claim 3 for causing the computer to perform the further  
2        functions of:

3              accepting a change of scale request from the user;  
4              in response to accepting said change of scale request, consolidating at least some of  
5              said plurality of icons with each other to form a second plurality of icons that is fewer in  
6              number than the number of icons in said first-mentioned plurality of icons; and

in response to accepting said change of scale request, redisplaying said domain using a decreased scale and also using the second plurality of icons to identify the locations of the records identified by said plurality of record identifiers.

1        14. The interface program of claim 3 for causing the computer to perform the further  
2        functions of:

3              accepting from the user a specification of an electronic note which has an associated  
4              location within the map; and

5              displaying a sticky-note icon on the map at position that corresponds to the associated  
6              location.

1        15. The interface program of claim 14 wherein the electronic note comprises a web  
2        page with its own externally accessible address which enables people to electronically access  
3        its contents through that address.

1        16. A database system stored on a computer-readable medium for causing a  
2        computer system to perform the functions of:

3              accepting search criteria including at least one of: (1) text; (2) a domain identifier  
4              identifying a domain; and (3) a filter identifier identifying a filter; and

5              retrieving a plurality of record identifiers each of which identifies a corresponding  
6              record which has associated therewith the text, domain identifier, or layer identifier of the  
7              search criteria, where the retrieving is performed with a spatial-keyword document index.

1        17. A method of retrieving a plurality of record identifiers each of which identifies a  
2        corresponding record which has associated therewith at least one of a text, a domain  
3        identifier, and a layer identifier as specified by search criteria, wherein such retrieving is  
4        performed with a spatial-keyword document index.

1        18. The method of claim 17 wherein the spatial-keyword document index comprises  
2        a spatial index tree extended to reference documents and a plurality of trees with the same  
3        structure as said spatial index tree but trimmed for particular lexicon entries and filters.

1           19. The method of claim 17 wherein said plurality of record identifiers are retrieved  
2       in a spatial-keyword document index tree and the branching structure of said tree is analyzed  
3       to identify geographic phenomena.

1           20. The method of claim 19 wherein a geographic phenomena is a spatial-keyword  
2       document index tree branching structure in which more than a predetermined fraction of the  
3       branches share more than predetermined number of parent nodes.

1           21. A program stored on a computer-readable medium for causing a computer system  
2       to perform the functions of:

3           loading documents referred to by document addresses;  
4           parsing those documents for more document addresses to load; and  
5           parsing those documents for possible spatial identifiers.

1           22. The program of claim 21 for causing the computer system to perform the further  
2       function of analyzing the possible spatial identifiers to determine a location in a domain.

1           23. The program of claim 21 wherein a portion of the document addresses are  
2       gathered by a metasearcher process, which queries other computer systems using text that  
3       refers to spatial domains.

1           24. The program of claim 21 for causing the computer system to perform the further  
2       function of computing a relevance score for each of the the possible spatial identifiers found  
3       in each document.

1           25. The program of claim 21 wherein the relevance score comprises one or more of:  
2           (1) the possible spatial identifier's position in the document;  
3           (2) the number of other possible spatial identifiers in the document;  
4           (3) whether the possible spatial identifier is in a sentence or is free standing; and  
5           (4) the formatted emphasis of the characters in the possible spatial identifier.

1           26. The program of claim 21 for causing the computer system to perform the further  
2       function of sorting the document addresses by relevance score before loading.

1           27. A program stored on a computer-readable medium for causing a computer system  
2       to perform the functions of:  
3           loading documents referred to by document addresses;  
4           parsing those documents for more document addresses to load;  
5           parsing those documents for possible spatial identifiers; and  
6           analyzing the possible spatial identifiers to determine a location in a domain.

1           28. A method of displaying spatially coded information, comprising:  
2           through an automated computer process, gathering documents in a database;  
3           selecting a subset of the documents which can be determined to contain spatial  
4       information;  
5           associating at least one spatial identifier with each document in the subset;  
6           indexing the documents, the indexing comprising an index on spatial identifiers and  
7       an index on keywords;  
8           providing a computer interface through which a user can submit a query comprising  
9       spatial information;  
10          responding to the query with a result set comprising documents; and  
11          displaying the result set to the user through the computer interface.

1           29. The method of claim 28, wherein the result set, when it contains more than one  
2       element, comprises a plurality of groups organized by spatial proximity, each group  
3       containing at least one document of the result set.

1           30. The method of claim 29 wherein the plurality of groups is ordered according to:  
2       predetermined function on groups representing relevance to the criteria.

1           31. The method of claim 29, wherein the content of each group is ordered according  
2       to a predetermined function on elements representing relevance to the criteria.

1           32. The method of claim 28, wherein the criteria include keywords.

1           33. A method for populating a spatial document database with hyperlinked  
2        documents containing spatial information, the method comprising:  
3           providing a destination database containing potential sources of gatherable  
4        documents;

5           providing a history database of known sources where documents have been gathered;  
6           providing a crawler computer process which can follow a hyperlink in a document to  
7        access a potential source of gatherable documents specified by the hyperlink;  
8           bootstrapping the crawler;  
9           iterating the crawler over the destination database, including the steps of:  
10          moving a potential source of gatherable documents from the destination database to  
11        the history database;  
12          inspecting the potential source for gatherable documents;  
13          storing any such gatherable documents in the spatial document database; and  
14          adding to the destination database all potential sources of gatherable documents  
15        which are referenced by a hyperlink in the gatherable documents.

1           34. The method of claim 33, wherein the bootstrapping comprises  
2           providing a plurality of locations of known interest;  
3           providing the destination database with a plurality of metasources, each metasource  
4        being a source of potential sources of gatherable documents, and each metasource responding  
5        to queries by the computer process with a result set comprising potential sources of  
6        gatherable documents; and  
7           priming the destination database by repeatedly running a primer process comprising:  
8           formulating a query based on at least one of the locations of known interest;  
9           directing the query to a metasource;  
10          receiving a result set for the query; and  
11          inserting each element of the result set into the destination database.

1           35. A method for scoring the relevance to a query of a document containing a known  
2 spatial identifier, the method comprising:

3           scoring the document for its relevance to a place specified in the query;

4           scoring the document for its relevance to a word specified in the query;


5           scoring the document for its quality; and

6           combining the scores to form a single score.

1           36. A method for indexing a plurality of documents to enable queries comprising  
2 keywords and spatial information, the method comprising:

3           initializing a master spatial tree of predetermined degree as a computer data structure,  
4 such that: (1) each leaf node of the master spatial tree represents a document; (2) each non-  
5 leaf node of the master spatial tree represents a range of space; (3) a root node of the master  
6 spatial tree represents a range of space encompassing any space that a valid query may refer  
7 to; and (4) the collection of non-leaf child nodes, relative to a parent node, define a partition  
8 on the range of space represented by the parent;

9           recursively adding a child node to the tree.



**FIG. 1**  
SUBSTITUTE SHEET (RULE 26)

217




FIG. 2

## **SUBSTITUTE SHEET (RULE 26)**

3/7




FIG. 3

4/7




FIG. 4

5/7




FIG. 5

6/7




FIG. 6

7/7




FIG. 7

SUBSTITUTE SHEET (RULE 26)

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US01/40173

**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(7) : G06F 17/30  
 US CL : 707/4, 104, 1-3, 5-8, 10; 345/348

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 707/4, 104, 1-3, 5-8, 10; 345/348

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
EAST**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,991,781 A (NIELSEN) 23 November 1999 (23.11.1999), ALL.	1-36
A	US 5,920,856 A (SYEDA-MAHMOOD) 06 July 1999 (06.07.1999), ALL.	1-36
A	US 5,802,361 A (WANG et al) 01 September 1998 (01.09.1998), ALL.	1-36

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier application or patent published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

20 June 2001 (20.06.2001)

Date of mailing of the international search report

28 JUN 2001

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks  
 Box PCT  
 Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

Uyen T Le

*Peggy Hanold*

Telephone No. 305-9000

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:



**BLACK BORDERS**



**IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**



**FADED TEXT OR DRAWING**




**BLURRED OR ILLEGIBLE TEXT OR DRAWING**



**SKEWED/SLANTED IMAGES**



**COLOR OR BLACK AND WHITE PHOTOGRAPHS**



**GRAY SCALE DOCUMENTS**



**LINES OR MARKS ON ORIGINAL DOCUMENT**



**REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**



**OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)